

ABSTRACT

A method and system for offloading I/O processing from a first computer to a second computer, using RDMA-capable network interconnects, are disclosed. The method and system include a client on the first computer communicating over an RDMA connection to a server on the second computer by way of a lightweight input/output (LWIO) protocol. The protocol generally comprises a network discovery phase followed by an I/O processing phase. During the discovery phase, the client and server determine a minimal list of shared RDMA-capable providers. During the I/O processing phase, the client posts I/O requests for offloading to the second machine over a mutually-authenticated RDMA channel. The I/O model is asymmetric, with read operations being implemented using RDMA and write operations being implemented using normal sends. Read and write requests may be completed in polling mode and in interrupt mode. Buffers are managed by way of a credit mechanism.